arXiv:1604.07994v1 [math.NA] 27 Apr 2016

Highly sparse surface couplings for subdomain-wise isoviscous Stokes finite element discretizations

Markus Huber[1] Ulrich Rüde[2] Christian Waluga[3]
Barbara Wohlmuth[1]

April 28, 2016

Abstract

The Stokes system with constant viscosity can be cast into different formulations by exploiting the incompressibility constraint. For instance the strain in the weak formulation can be replaced by the gradient to decouple the velocity components in the different coordinate directions. Thus the discretization of the simplified problem leads to fewer nonzero entries in the stiffness matrix. This is of particular interest in large scale simulations where a reduced memory bandwidth requirement can help to significantly accelerate the computations. In the case of a piecewise constant viscosity, as it typically arises in multi-phase flows, or when the boundary conditions involve traction, the situation is more complex, and one has to treat the cross derivatives in the original Stokes system with care. A naive application of the standard vectorial Laplacian results in a physically incorrect solution, while formulations based on the strain increase the computational effort everywhere, even when the inconsistencies arise only from an incorrect treatment in a small fraction of the computational domain. Here we propose a new approach that is consistent with the strain-based formulation and preserves the decoupling advantages of the gradient-based formulation in isoviscous subdomains. The modification is equivalent to locally changing the discretization stencils, hence the more expensive discretization is restricted to a lower dimensional interface, making the additional computational cost asymptotically negligible. We demonstrate the consistency and convergence properties of the method and show that in a massively parallel setup, the multigrid solution of the resulting discrete systems is faster than for the classical strain-based formulation. Moreover, we give an application example which is inspired by geophysical research.

[1]Institute for Numerical Mathematics (M2), Technische Universität München, Boltzmannstrasse 3, D–85748 Garching b. München, Germany
[2]Department of Computer Science 10, FAU Erlangen-Nürnberg, Cauerstraße 6, D–91058 Erlangen, Germany
[3]liNear GmbH, Kackertstr. 11, D–52072 Aachen, Germany; corresponding author

1 Introduction

The non-isoviscous Stokes equations are an important mathematical model for a wide range of applications and serve as a building block for more complicated fluid models. In several applications of interest, however, the model can be considerably simplified by formally decomposing the simulation domain into isoviscous sub-domains. To this end, consider the incompressible Stokes equations

$$-\operatorname{div} \boldsymbol{\sigma}(\boldsymbol{u}, p) = \boldsymbol{f}, \tag{1a}$$

$$\operatorname{div} \boldsymbol{u} = 0, \tag{1b}$$

in a polyhedral domain $\Omega \subset \mathbb{R}^d$, $d = 2, 3$, where $\boldsymbol{u} = [u_1, \dots, u_d]^\top$ denotes the velocity, p the pressure, and $\boldsymbol{f} = [f_1, \dots, f_d]^\top]$ is an external forcing term. The stress tensor is defined as

$$\boldsymbol{\sigma}(\boldsymbol{u}, p) := 2\mu \operatorname{sym} \nabla \boldsymbol{u} - p \cdot \mathrm{Id}, \tag{2}$$

where $\operatorname{sym} T := \frac{1}{2}(T + T^\top)$ denotes the symmetric part of a tensor $T \in \mathbb{R}^{d\times d}$, and μ is a positive scalar viscosity field. The focus of this work lies on a model setting resulting from the assumption that the viscosity is a piecewise constant function with a possibly large contrast, i.e.,

$$\mu = \begin{cases} \mu_1 & \text{in } \Omega_1, \\ \mu_2 & \text{in } \Omega_2 \end{cases} \tag{3}$$

with $\mu_1 \geq \mu_2 > 0$.

Such models have relevant applications, for instance in incompressible two-phase flow simulations (cf. [16] and the references therein), or in fundamental geophysical studies; cf. e.g. [17, 9, 39].

To simplify the exposition, the disjoint subdomains Ω_1 and Ω_2 are assumed to be polyhedral such that $\overline{\Omega} = \overline{\Omega}_1 \cup \overline{\Omega}_2$, and the interface between both subdomains is denoted by $\Gamma_{12} := \partial\Omega_1 \cap \partial\Omega_2$. The equations (1) are then equivalent to the interface formulation

$$-\mu_i \Delta \boldsymbol{u} + \nabla p = \boldsymbol{f}, \qquad \text{in } \Omega_i,\ i = 1, 2 \tag{4a}$$

$$\operatorname{div} \boldsymbol{u} = 0, \qquad \text{in } \Omega_i, \tag{4b}$$

$$[\![\boldsymbol{u}]\!] = 0, \quad [\![\boldsymbol{\sigma}(\boldsymbol{u}, p) \cdot \boldsymbol{n}]\!] = 0, \qquad \text{on } \Gamma_{12}, \tag{4c}$$

which is discussed e.g. in [33, 22, 29]. Here $\boldsymbol{n}_i$ stands for the outward pointing unit normal on $\partial\Omega_i$, and we omit the index where there is no ambiguity. Moreover, $[\![v]\!] := v|_{\Gamma_{12}\cap\Omega_1} - v|_{\Gamma_{12}\cap\Omega_2}$ denotes the oriented jump of a scalar or vectorial quantity v across the interface, and $[\![T \cdot \boldsymbol{n}]\!] := T|_{\Gamma_{12}\cap\Omega_1} \cdot \boldsymbol{n}_1 + T|_{\Gamma_{12}\cap\Omega_2} \cdot \boldsymbol{n}_2$ stands for the jump of the normal component of a tensorial quantity T.

In terms of boundary conditions, we complement the equations (4) by *homogeneous Dirichlet* conditions $\boldsymbol{u} = \boldsymbol{0}$ on $\Gamma_D \subset \partial\Omega$ and prescribe either *traction-free* boundary conditions

$$\boldsymbol{\sigma}(\boldsymbol{u}, p) \cdot \boldsymbol{n} = \boldsymbol{0}, \qquad \text{on } \Gamma_F \subset \partial\Omega, \tag{5a}$$

or *free-slip* boundary conditions

$$\boldsymbol{u}\cdot\boldsymbol{n}=0,\quad \boldsymbol{n}\times\boldsymbol{\sigma}(\boldsymbol{u},p)\cdot\boldsymbol{n}=0,\qquad \text{on } \Gamma_F\subset\partial\Omega. \tag{5b}$$

As it is standard, we assume $\Gamma_D\cap\Gamma_F=\emptyset$, $\bar{\Gamma}_D\cup\bar{\Gamma}_F=\partial\Omega$, and the $(d-1)$-dimensional measure of Γ_D is positive such that Korn-type estimates hold.

The important difference between the single domain (1) and the multi-domain formulation (4) is that the presence of cross-derivatives in (4) is limited to Γ_{12}. At the interface Γ_{12}, the continuity of the normal stresses is enforced via (4c). Moreover at the boundary part Γ_F, the traction is (partly) constrained via (5). Thus in the following, both Γ_F and Γ_{12} will be called *interface* to simplify notation. The equivalence of (1) and (4) results from the fact that for solenoidal velocities, we have

$$2\,\mathrm{div}(\mathrm{sym}\,\nabla\boldsymbol{u})=\Delta\boldsymbol{u}+\nabla\,\mathrm{div}\,\boldsymbol{u}=\Delta\boldsymbol{u}, \tag{6}$$

which decouples the velocity components in the momentum equation. Preserving this decoupling property in a discrete sense is of special interest in large-scale computations since the number of nodes at the interface increases with $\mathcal{O}(h^{-d+1})$ while the total number of degrees of freedom grows with $\mathcal{O}(h^{-d})$. Thus, asymptotically, the number of nodes associated with the interior of the subdomains dominates, and the computational cost can be significantly reduced by using an approach that decouples in the interior of the isoviscous subdomains. This observation does not only hold for traditional element-wise assembling procedures but also for the matrix-vector product with the stiffness matrix. A fast realization of these basic operations is essential to obtain a fast time-to-solution of many iterative solvers such as multigrid methods. It is however not obvious how to exploit the previously discussed decoupling property for large-scale applications. This fact is related to the observation that interface problems often result in a saddle point structure that may give rise to stability concerns.

To elaborate on this, let us choose a natural setting for the interface problem by defining appropriate function spaces. For the velocity space we choose $\boldsymbol{V}:=\{\boldsymbol{v}\in H^1(\Omega)^d\ :\ \boldsymbol{v}=\boldsymbol{0}\text{ on }\Gamma_D,\int_{\Gamma_F}\boldsymbol{v}\cdot\boldsymbol{n}\,ds=0\}$ and $\boldsymbol{V}:=\{\boldsymbol{v}\in H^1(\Omega)^d\ :\ \boldsymbol{v}=\boldsymbol{0}\text{ on }\Gamma_D,\boldsymbol{v}\cdot\boldsymbol{n}=0\text{ on }\Gamma_F\}$ for (5a) and (5b), respectively. We point out that the definition of the velocity space automatically guarantees that $\int_{\partial\Omega}\boldsymbol{v}\cdot\boldsymbol{n}\,ds=\int_\Omega\mathrm{div}\,\boldsymbol{v}\,dx=0$, $\boldsymbol{v}\in\boldsymbol{V}$ and thus for all our settings, we only work with $Q:=\{\,q\in L^2(\Omega)\ :\ \int_\Omega\mu^{-1}q\,\mathrm{d}x=0\,\}$ as test and solution space for the pressure. This choice ensures uniqueness of the pressure with respect to the constant mode. The weak problem will then read as: given $\boldsymbol{f}\in\boldsymbol{V}'$, find $(\boldsymbol{u},p)\in\boldsymbol{V}\times Q$, such that

$$a(\boldsymbol{u},\boldsymbol{v})+b(\boldsymbol{v},p)=\langle\boldsymbol{f},\boldsymbol{v}\rangle,\qquad \text{for all }\boldsymbol{v}\in\boldsymbol{V}, \tag{7a}$$

$$b(\boldsymbol{u},q)=0,\qquad \text{for all } q\in Q, \tag{7b}$$

where $b(\boldsymbol{v},q):=-(\mathrm{div}\,\boldsymbol{v},q)_0$ is the standard weak form of the divergence and $(\cdot,\cdot)_0$ denotes the L^2 scalar product on Ω. The exact form of $a(\cdot,\cdot)$ however

depends on the strong formulation on which we base our derivation. Starting from (1) leads to the bilinear form

$$a_{\mathrm{sym}}(\boldsymbol{u},\boldsymbol{v}) := (2\mu \operatorname{sym}\nabla\boldsymbol{u}, \operatorname{sym}\nabla\boldsymbol{v}). \tag{8}$$

By Korn's second inequality, we obtain positive definiteness of $a_{\mathrm{sym}}(\cdot,\cdot)$ on $\boldsymbol{V}\times\boldsymbol{V}$. The drawback of this weak form is that it does not exploit the knowledge that the velocity components are decoupled in the strong momentum balance inside the subdomains, which leads to a reduced sparsity in the finite-dimensional setting.

Another straightforward approach would be to start from the interface formulation (4). Integrating by parts over each subdomain and incorporating the momentum balance, we obtain a weak Laplacian which is augmented with consistency terms at the interface

$$a_{\mathrm{int}}(\boldsymbol{u},\boldsymbol{v}) := (\mu\nabla\boldsymbol{u},\nabla\boldsymbol{v}) + \sum_{i=1}^{2}(\mu\nabla\boldsymbol{u}^{\top}\boldsymbol{n},\boldsymbol{v})_{\partial\Omega_i\cap\Gamma_{12}} + (\mu\nabla\boldsymbol{u}^{\top}\boldsymbol{n},\boldsymbol{v})_{\Gamma_F}, \tag{9}$$

cf. e.g. [23], where a similar construction for the special case of boundary conditions involving the surface traction is considered. The above formulation, while being consistent and decoupled inside the subdomains, has two major drawbacks: Firstly, due to the additional consistency terms, this form sacrifices symmetry. Secondly, it is – to the authors' best knowledge – not possible to prove coercivity of $a_{\mathrm{int}}(\cdot,\cdot)$ for large viscosity contrasts. Let us at this point underline that simply neglecting the consistency terms in a discrete solver will, in general, result in convergence to a wrong solution if $\mu_1 \neq \mu_2$ or $\Gamma_F \neq \emptyset$. Limache et al. [23] interpret the lack of consistency as a violation of the principle of objectivity, a basic axiom of continuum mechanics.

In this article, we introduce a framework for incompressible flow with subdomain-wise constant viscosities, that exploits the increased sparsity of the vectorial Laplacian and thus – assuming comparable convergence rates – allows for a significantly faster iterative solution as compared to the classical strain based formulation. Our goal is to provide a mathematical sound reformulation in a variationally consistent setting and to report on the resulting performance and scalability results. The paper is organized as follows: In the following Section 2, we shall derive an alternative bilinear form which is both consistent and symmetric, therefore eliminates the major drawback of the form $a_{\mathrm{int}}(\cdot,\cdot)$ defined above. Subsequently, in Section 3 we discuss the circumstances under which the modified form leads us to a stable and convergent approximation scheme in case of finite dimensional velocity and pressure spaces. Particularly, we present and analyze a lowest-order stabilized method which has a decoupling property built-in and suppresses spurious pressure modes. In Section 4, we give some illustrative numerical examples to complement our theory and to report on the discretization error for the different formulations. Section 5 is devoted to performance considerations. More precisely, we consider the influence of the formulation on the time to solution for a highly scalable and parallel all-at-once

multigrid solver. Finally in Section 6, we present a simplified geophysical model example and show that the new modified scheme yields results which are in agreement with the physically accepted strain-based formulation of the Stokes problem.

2 A symmetric decoupled weak formulation

In this section, we provide two different weak formulations that decouple in the subdomains and involve cross derivative terms only across the interface. The approach proposed in this work is primarily based on an observation that can be traced back at least to Pironneau [31, Sect. 6.4], namely, that we can rewrite

$$(2\operatorname{sym}\nabla\boldsymbol{u},\operatorname{sym}\nabla\boldsymbol{v})_{\Omega_i}-(\operatorname{div}\boldsymbol{u},\operatorname{div}\boldsymbol{v})_{\Omega_i}=(\nabla\boldsymbol{u},\nabla\boldsymbol{v})_{\Omega_i}+(\nabla\boldsymbol{u}^\top\boldsymbol{n}-\operatorname{div}\boldsymbol{u}\cdot\boldsymbol{n},\boldsymbol{v})_{\partial\Omega_i}. \tag{10}$$

For a formal proof and suitable assumptions, we refer to Lemma 2 below. Using (10), we can easily derive a symmetric bilinear form

$$\begin{aligned} a_{\mathrm{tr}}(\boldsymbol{v},\boldsymbol{u}) &:= a_{\mathrm{sym}}(\boldsymbol{v},\boldsymbol{u})-(\mu\operatorname{div}\boldsymbol{u},\operatorname{div}\boldsymbol{v}) \\ &= a_{\mathrm{dev}}(\boldsymbol{u},\boldsymbol{v})-(\mu(1-\tfrac{2}{d})\operatorname{div}\boldsymbol{u},\operatorname{div}\boldsymbol{v}), \end{aligned} \tag{11}$$

where $a_{\mathrm{dev}}(\boldsymbol{u},\boldsymbol{v}) := (2\mu\operatorname{dev}\operatorname{sym}\nabla\boldsymbol{u},\operatorname{dev}\operatorname{sym}\nabla\boldsymbol{v})$. Here $\operatorname{dev}T := T-\frac{1}{d}\operatorname{tr}(T)\cdot\mathrm{Id}$ denotes the deviatoric part of a tensor T. Thus, we directly obtain a decoupling property for the momentum part also in the weak setting. The above modification amounts to formally changing the stress tensor to

$$\begin{aligned} \boldsymbol{\sigma}(\boldsymbol{u},p) &= 2\mu\operatorname{sym}\nabla\boldsymbol{u}-(p+\mu\operatorname{div}\boldsymbol{u})\cdot\mathrm{Id} \\ &= 2\mu\operatorname{dev}\operatorname{sym}\nabla\boldsymbol{u}-(p+\mu(1-\tfrac{2}{d})\operatorname{div}\boldsymbol{u})\cdot\mathrm{Id}\,. \end{aligned} \tag{12}$$

It is obvious that $a_{\mathrm{tr}}(\cdot,\cdot)$ is symmetric, and moreover for $d=2$ it is also coercive on $\boldsymbol{V}\times\boldsymbol{V}$ if $\Gamma_D=\partial\Omega$ with respect to the standard H^1-norm, see, e.g., [27, Lemma 4.2]. For solenoidal elements of $\boldsymbol{V}$ the bilinear form (11) is equal to (8), and thus the modified problem is strongly consistent with the original Stokes problem (1) due to the incompressibility constraint $\operatorname{div}\boldsymbol{u}=0$.

Before we discuss finite element approximations for the corresponding weak problem, we shall discuss two possible equivalent reformulations which can be used for the implementation to exploit the advantageous decoupling property.

2.1 Equivalence with a tangential interface coupling

Before we can state our main results, we need to discuss some preliminary algebraic identities. These identities provide a useful tool to expand a trace-free tensor in terms of a rotational basis. For the two-dimensional setting, $d=2$, we define the orthogonal matrix

$$R_2^1=\begin{pmatrix}0&-1\\1&0\end{pmatrix}$$

and set $n_d = n_2 = 1$. For the three-dimensional case, $d = 3$, we accordingly define

$$R_3^1 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}, \qquad R_3^2 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \qquad R_3^3 = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

and set $n_d = n_3 = 3$. For $\mathbf{x} \in \mathbb{R}^2$ we have $R_2^1 \mathbf{x} \perp \mathbf{x}$ and for $\mathbf{x} \in \mathbb{R}^3$ we get $R_3^j \mathbf{x} = \mathbf{e}_j \times \mathbf{x}$, where $\{\mathbf{e}_j, 1 \le j \le 3\}$, denotes the canonical basis of $\mathbb{R}^3$.

Lemma 1. *Let $T \in \mathbb{R}^{d\times d}$, $d = 2, 3$. Then we have the equalities*

$$T - \operatorname{tr}(T) \cdot \mathrm{Id} = \sum_{j=1}^{n_d} R_d^j T^\top R_d^j, \tag{13}$$

$$\operatorname{dev} T = \sum_{j=1}^{n_d} R_d^j \operatorname{dev} T^\top R_d^j. \tag{14}$$

Proof. Both equalities are easily obtained from a straightforward computation. For $d = 3$, it is helpful to note that $\sum_{j=1}^3 R_d^j R_d^j = -2\,\mathrm{Id}$, and

$$\begin{aligned} \operatorname{dev} T = T - \operatorname{tr}(T) \cdot \mathrm{Id} + \frac{2}{3}\operatorname{tr}(T) \cdot \mathrm{Id} &= \sum_{j=1}^3 R_d^j T^\top R_d^j - \frac{1}{3}\operatorname{tr}(T) \sum_{j=1}^3 R_d^j R_d^j \\ = \sum_{j=1}^3 R_d^j (T^\top - \frac{1}{3}\operatorname{tr}(T^\top) \cdot \mathrm{Id}) R_d^j &= \sum_{j=1}^3 R_d^j \operatorname{dev} T^\top R_d^j, \end{aligned}$$

thus (14) follows from (13) and the definition of the deviatoric operator. □

The previous identities are essential for our derivation as we will show next. Let $\mathcal{T}_h(\Omega_i)$ be a family of shape regular triangulations on Ω_i such that each vertex on Γ_{12} is a vertex of both $\mathcal{T}_h(\Omega_1)$ and $\mathcal{T}_h(\Omega_2)$, with at least one inner vertex on Γ_{12}.

Lemma 2. *For all $\boldsymbol{u}, \boldsymbol{v} \in \boldsymbol{H}^1(\Omega_i) \cap \prod_{T\in\mathcal{T}_h(\Omega_i)} \boldsymbol{H}^2(T)$ there holds:*

$$(\nabla \boldsymbol{u}^\top - \operatorname{div} \boldsymbol{u} \cdot \mathrm{Id}, \nabla \boldsymbol{v})_{\Omega_i} = (\nabla \boldsymbol{u}^\top \boldsymbol{n} - \operatorname{div} \boldsymbol{u} \cdot \boldsymbol{n}, \boldsymbol{v})_{\partial\Omega_i}. \tag{15}$$

Moreover, we can equivalently rewrite the surface terms by tangential derivatives as

$$(\nabla \boldsymbol{u}^\top \boldsymbol{n} - \operatorname{div} \boldsymbol{u} \cdot \boldsymbol{n}, \boldsymbol{v})_{\partial\Omega_i} = -\sum_{j=1}^{n_d} (\nabla \boldsymbol{u}\, R_d^j \boldsymbol{n}, R_d^j \boldsymbol{v})_{\partial\Omega_i}. \tag{16}$$

Proof. The identity (15), which readily implies (10), can be found in [31, Sect. 6.4; eq. (130)]. Since the proof is omitted there, we shall give one here for the sake of completeness: We integrate by parts on each element inside Ω_i and notice that,

due to the regularity requirements, we have $\nabla \mathrm{div}\, \boldsymbol{u} - \mathrm{div}\, \nabla \boldsymbol{u}^\top = \boldsymbol{0}$ in every $T \in \mathcal{T}_h(\Omega_i)$. By this, we find after summation that

$$(\nabla \boldsymbol{u}^\top - \mathrm{div}\, \boldsymbol{u} \cdot \mathrm{Id}, \nabla \boldsymbol{v})_{\Omega_i} = \sum_{T \in \mathcal{T}_h(\Omega_i)} (\nabla \boldsymbol{u}^\top \boldsymbol{n} - \mathrm{div}\, \boldsymbol{u} \cdot \boldsymbol{n}, \boldsymbol{v})_{\partial T}.$$

Now it is important to see that $\nabla \boldsymbol{u}^\top \cdot \boldsymbol{n} - \mathrm{div}\, \boldsymbol{u} \cdot \boldsymbol{n}$ is a derivative acting only in tangential direction, and hence the contributions of each facet in the interior of Ω_i cancel out due to the inverse orientation of their normals. To show this, we make use of the identity (13) with $T = \nabla \boldsymbol{u}^\top$, which directly yields the second identity as a side product. □

Given the above considerations and defining the tangential vectors $\boldsymbol{t}_j := R_d^j \boldsymbol{n}$, $j = 1, \ldots, n_d$, and $\boldsymbol{v}_j^\perp := R_d^j \boldsymbol{v}$, it is easy to see that on $\boldsymbol{V} \times \boldsymbol{V}$ the bilinear form $a_{\mathrm{tr}}(\boldsymbol{u}, \boldsymbol{v})$ can be equivalently written as

$$a_{\mathrm{tr}}(\boldsymbol{u}, \boldsymbol{v}) = (\mu \nabla \boldsymbol{u}, \nabla \boldsymbol{v}) - [\![\mu]\!] \sum_{j=1}^{n_d} (\partial_{\boldsymbol{t}_j} \boldsymbol{u}, \boldsymbol{v}_j^\perp)_{\Gamma_{12}} - \mu \sum_{j=1}^{n_d} (\partial_{\boldsymbol{t}_j} \boldsymbol{u}, \boldsymbol{v}_j^\perp)_{\Gamma_F}. \tag{17}$$

This formulation allows us to see that the additional couplings at the interface are in fact only introduced in tangential direction, i.e, in conforming velocity spaces, the cross-couplings are restricted to the coefficients at the interface.

2.2 Equivalence with a modified interface stencil

Although the representation (17) allows us to avoid cross derivative terms within the subdomain, it requires the assembly of interface terms which may not be straightforwardly supported in existing codes. We shall thus shed some light on an alternative interpretation of the interface treatment.

In massively parallel large-scale simulations, it is becoming increasingly important to develop matrix-free implementations since they do not only save the memory for storing the full operator but also the effort to retrieve the matrix entries in each iteration from memory which is often the dominating cost factor. In particular, matrix-free implementations are of interest for hierarchically refined grids that result in patch-wise uniform fine meshes; cf. e.g. [2, 12, 13] and the references contained therein. Here, assuming constant coefficients, the interior stencils can be efficiently pre-computed for the coarse mesh elements and can be applied on the fly on the fine mesh. A modified approach is thus in our context only necessary at the interface, where we propose to express the contributions of the bilinear form $a_{\mathrm{tr}}(\cdot, \cdot)$ in a stencil-based fashion.

Let us consider a finite-dimensional space $\mathbf{V}_h \subset \mathbf{V}$. Moreover let us assume that we can decompose $\boldsymbol{v}_h = \boldsymbol{v}_h^\circ + \widehat{\boldsymbol{v}}_h$ such that $\boldsymbol{v}_h^\circ \in \boldsymbol{V}_h$ and $\boldsymbol{v}_h^\circ|_{\Omega_i} \in \boldsymbol{H}_0^1(\Omega_i)$, $i = 1, 2$. It is easy to see that we obtain

$$a_{\mathrm{tr}}(\boldsymbol{u}_h, \boldsymbol{v}_h) = (\mu \nabla \boldsymbol{u}_h, \nabla \boldsymbol{v}_h^\circ)_\Omega + a_{\mathrm{tr}}(\boldsymbol{u}_h, \widehat{\boldsymbol{v}}_h), \tag{18}$$

where we use (11) for the evaluation of $a_{\mathrm{tr}}(\boldsymbol{u}_h, \widehat{\boldsymbol{v}}_h)$.

The conceptual sketch in Fig. 1 illustrates the coupling of the x- and y-velocity components in two dimensions. The standard bilinear form (8) depicted in the second picture from the left results in a fully coupled setting, while the bilinear form (11) yields a reduced coupling in the components. We note that a closer look reveals that both stiffness matrices resulting from (17) and (18) are mathematically equal, but differ in their sparsity structure in a stencil based code which is illustrated by the translucent/green box.

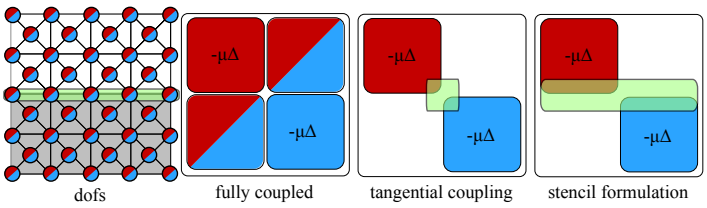

Figure 1: An illustration of different types of cross derivative couplings.

Using (17), we obtain a stiffness matrix where only the nodes at the interface couple in x and y direction, see the second picture from the right. Using the equivalent form (18) results in the picture on the right where formally also coupling terms between the subdomain velocity components and interface test functions occur. Although, the element-wise assembled contributions cancel out within the global assembling process, the stencil entry is formally filled by a zero and contributes to the floating point operations (FLOPS) as well as the memory footprint of the operator application. However, since the additional entries are only relevant at the interface, these extra costs can be asymptotically neglected. From the point of view of stencil-based codes, the application of the viscous operator can thus be realized by applying the Laplacian stencils inside the domains of same viscosity and the full stencil only at the interface Γ_{12} between the subdomains with different viscosity and on Γ_F to ensure consistency. Hence, this decoupling of the velocities inside each subdomain allows for an efficient matrix vector multiplication as required for many smoothers in the multigrid context, since each velocity component can be treated independently.

Remark 1. *The decoupling property is not restricted to settings in which the mesh resolves the interface. For instance, in two-phase flow computations, where the interface is described by a level-set function, the same techniques can be used. In this case, we have to differentiate between shape-functions which are supported at the interface and those which are not. Given such a decomposition, the decoupling mechanisms described above can be straightforwardly applied to increase the sparsity of the global stiffness matrix. The computational effort of identifying the corresponding interface degrees of freedom will anyway be needed for the element based assembly routines so that this technique can be directly incorporated without any significant extra cost. The decoupling property also generalizes to higher order finite elements.*

3 Stability and uniform coercivity in the discrete setting

By construction, the saddle point problem (7) with $a(\cdot,\cdot)$ given by $a_{\mathrm{tr}}(\cdot,\cdot)$ is consistent. We associate with $(\boldsymbol{v},q)\in \boldsymbol{V}\times Q$ the norm

$$\begin{aligned}\|(\boldsymbol{v},q)\|^2 &:= \|\boldsymbol{v}\|_{\boldsymbol{V}}^2 + \|q\|_Q^2,\\ \|\boldsymbol{v}\|_{\boldsymbol{V}}^2 &:= a_{\mathrm{sym}}(\boldsymbol{v},\boldsymbol{v}), \qquad \|q\|_Q^2 := (q,q)_Q, \quad (p,q)_Q := (p/(2\mu),q)_0\end{aligned}$$

and consider additionally on $[H^1(\Omega)]^d$ the semi-norm $|\boldsymbol{v}|_{\boldsymbol{V}}^2 := a_{\mathrm{dev}}(\boldsymbol{v},\boldsymbol{v})$.

Remark 2. *The choice of our boundary conditions and Korn's inequality guarantee that $\|\cdot\|_{\boldsymbol{V}}$ is equivalent to the standard H^1-Hilbert space norm on $\boldsymbol{V}$. However, in general, the constants in the bounds depend on μ. If on Ω_1, i.e., on the subdomain with the larger viscosity, a non-trivial Dirichlet boundary part is prescribed, then the norm $\|\cdot\|_{\boldsymbol{V}}$ is equivalent to $(2\mu\nabla\boldsymbol{v},\nabla\boldsymbol{v})_0^{0.5}$ with constants independent of μ. The inf-sup condition on $\boldsymbol{V}\times Q$ is shown in [29] with respect to the norm pairing $(2\mu\nabla\cdot,\nabla\cdot)_0^{0.5}$ and $\|\cdot\|_Q$, and thus it automatically holds true for the pairing $\|\cdot\|_{\boldsymbol{V}}$ and $\|\cdot\|_Q$.*

Introducing the subspace of solenoidal velocities $\boldsymbol{V}_0 := \{\boldsymbol{v}\in\boldsymbol{V} \ : \ \operatorname{div}\boldsymbol{v} = 0\}$, it is trivial to see that

$$a_{\mathrm{sym}}(\boldsymbol{v},\boldsymbol{v}) = a_{\mathrm{dev}}(\boldsymbol{v},\boldsymbol{v}) = a_{\mathrm{tr}}(\boldsymbol{v},\boldsymbol{v}), \quad \boldsymbol{v}\in\boldsymbol{V}_0,$$

and consequently the saddle point problem (7) with $a(\cdot,\cdot)$ given by $a_{\mathrm{tr}}(\cdot,\cdot)$ is stable in the continuous setting. However, the discrete setting does not necessarily inherit this property.

Before we consider stable discretizations of the saddle point problem, we comment on the semi-norm $|\cdot|_{\boldsymbol{V}}$. According to [27, Lemma 4.2], $|\cdot|_{\boldsymbol{V}}$ is equivalent to the standard H^1-Hilbert space norm on $\boldsymbol{V}$ for $d=3$. Moreover the equivalence to $(2\mu\nabla\boldsymbol{v},\nabla\boldsymbol{v})_0^{0.5}$ is uniform with respect to μ if on Ω_1 a non-trivial Dirichlet boundary part is prescribed. However, this does not hold for $d=2$. In the planar setting, the equivalence of $|\cdot|_{\boldsymbol{V}}$ to the standard H^1-Hilbert space norm on $\boldsymbol{V}$ is only guaranteed if $\Gamma_D=\partial\Omega$. This difference in the equivalence is related to the fact that the kernel $\mathcal{N}$ of the operator $\operatorname{dev}\operatorname{sym}\nabla$ is infinite dimensional for $d=2$ but finite dimensional for $d=3$. We refer to [32] for a sequence serving as counterexample in $d=2$. For $d=3$, the dimension of $\mathcal{N}$ equals 10, including the six dimensional space of rigid body modes. Since $\operatorname{div}\mathcal{N}\neq\{0\}$, we can find a $\boldsymbol{v}\in\mathcal{N}$ such that $a_{\mathrm{tr}}(\boldsymbol{v},\boldsymbol{v})<0$. Let us consider as example $\boldsymbol{v}:=\mathbf{x}$: we then get $a_{\mathrm{dev}}(\boldsymbol{v},\boldsymbol{v})=0$ and $a_{\mathrm{tr}}(\boldsymbol{v},\boldsymbol{v}) = -3(\mu_1|\Omega_1|+\mu_2|\Omega_2|)<0$. In contrast to $a_{\mathrm{dev}}(\cdot,\cdot)$, the bilinear form $a_{\mathrm{tr}}(\cdot,\cdot)$ is not positive semi-definite on $[H^1(\Omega)]^3$ and does not define a semi-norm.

In the light of the previous considerations, it becomes obvious that the stability of the previously discussed modified schemes is not straightforward to establish if the incompressibility constraint cannot be strongly satisfied given the combination of finite element spaces $\mathbf{V}_h\times Q_h\subset\mathbf{V}\times Q$ at hand. Such

mixed spaces, where $\operatorname{div} \mathbf{V}_h = Q_h$ are not very wide-spread in practice since they require higher-order ansatz functions and/or special meshes, cf. e.g. the discussions in [37, 40, 6, 8]. In particular, for the case $d = 3$, the stability analysis for the so-called Scott–Vogelius spaces has been an open problem for almost three decades; cf. the discussion in [28] for recent progress.

In the following we shall thus demonstrate techniques that allow to prove stability of our modified scheme for stabilized finite element discretizations.

3.1 Stabilized lowest-order discretization

Let us recall that the computational domain Ω is subdivided into a conforming simplicial mesh $\mathcal{T}_h$ that resolves the interface. For the velocity discretization, we use linear, conforming finite elements and for the pressure space we employ piecewise constants, i.e.,

$$\begin{aligned} \boldsymbol{V}_h &:= \{\boldsymbol{v} \in \boldsymbol{V} : \boldsymbol{v}|_T \in [P_1(T)]^d,\ \forall T \in \mathcal{T}_h\}, \\ Q_h &:= \{q \in Q : q|_T \in P_0(T),\ \forall T \in \mathcal{T}_h\}. \end{aligned}$$

In terms of a suitable stabilization bilinear form $c(\cdot,\cdot)$, the discrete weak problem then reads as: given $\boldsymbol{f} \in \boldsymbol{V}'$, find $(\boldsymbol{u}_h, p_h) \in \boldsymbol{V}_h \times Q_h$, such that

$$B(\boldsymbol{u}_h, p_h; \boldsymbol{v}_h, q_h) = \langle \boldsymbol{f}, \boldsymbol{v}_h \rangle, \qquad \text{for all } (\boldsymbol{v}_h, q_h) \in \boldsymbol{V}_h \times Q_h, \tag{19a}$$

$$B(\boldsymbol{u}_h, p_h; \boldsymbol{v}_h, q_h) := a_{\mathrm{tr}}(\boldsymbol{u}_h, \boldsymbol{v}_h) + b(\boldsymbol{v}_h, p_h) + b(\boldsymbol{u}_h, q_h) - c(p_h, q_h). \tag{19b}$$

The following stability analysis is based on a continuous inf-sup condition for the model problem [29, Thm. 1] and the use of a subdomain-wise defined Scott-Zhang type operator S_h preserving homogeneous boundary conditions [36]. To obtain a globally conforming approximation, we also use a suitable boundary modification for S_h at the interface Γ_{12}, i.e., the nodal values of $S_h\boldsymbol{v}$ at each node p on $\bar{\Gamma}_{12}$ are purely defined in terms of $\boldsymbol{v}|_{\bar{\Gamma}_{12}}$.

Theorem 1. *Let the stabilization term $c(\cdot,\cdot)$ satisfy*

$$c(q_h, q_h) \le \gamma_0 \|q_h\|_Q^2, \quad (q_h, \operatorname{div}(\boldsymbol{v} - S_h\boldsymbol{v}))_0^2 \le C_{stab}^2 c(q_h, q_h) \|\sqrt{2\mu}\nabla \boldsymbol{v}\|_0^2$$

with $C_{stab} < \infty$ and $\gamma_0 < \frac{2d}{d-2}$ for all $q_h \in Q_h$ and $\boldsymbol{v}|_{\Omega_i} \in [H_0^1(\Omega_i)]^d$. Then the variational formulation (19) is uniformly stable with respect to the norm $\|(\boldsymbol{v}, q)\|$ and has a unique solution. Moreover the finite element error can be bounded by the best approximation error, i.e.,

$$\|(\boldsymbol{u} - \boldsymbol{u}_h, p - p_h)\| \le C \inf_{(\boldsymbol{v},q) \in \boldsymbol{V}_h \times Q_h} \|(\boldsymbol{u} - \boldsymbol{v}, p - q)\|.$$

Proof. It is sufficient to show that there exists a $(\boldsymbol{v}_h, q_h) \in \boldsymbol{V}_h \times Q_h$ such that

$$\frac{B(\boldsymbol{u}_h, p_h; \boldsymbol{v}_h, q_h)}{\|(\boldsymbol{v}_h, p_h)\|} \ge c\|(\boldsymbol{u}_h, p_h)\|. \tag{20}$$

We recall that μ is piecewise constant and moreover $(2\mu \operatorname{div} \boldsymbol{u}_h, 1)_Q = (\operatorname{div} \boldsymbol{u}_h, 1)_0 = 0$ due to the definition of $\boldsymbol{V}$. Following [29, Thm. 1], there exists $\boldsymbol{z} \in \boldsymbol{V}$ such that

$$(\operatorname{div} \boldsymbol{z}, p_h + (1 - \frac{2}{d})\mu \operatorname{div} \boldsymbol{u}_h)_0 = -\|p_h + (1 - \frac{2}{d})\mu \operatorname{div} \boldsymbol{u}_h\|_Q^2 \tag{21a}$$

$$\|\sqrt{2\mu}\nabla \boldsymbol{z}\|_0 \leq C_1 \|p_h + (1 - \frac{2}{d})\mu \operatorname{div} \boldsymbol{u}_h\|_Q \tag{21b}$$

with a constant $C_1 < \infty$ independent of $(\boldsymbol{u}_h, p_h)$. Using the assumption on the triangulation, a closer look reveals that $\boldsymbol{z}$ can assume the special form $\boldsymbol{z} = \boldsymbol{z}_h + \boldsymbol{z}_0$ with $\boldsymbol{z}_h \in \boldsymbol{V}_h$, $\boldsymbol{z}_0|_{\Omega_i} \in [H_0^1(\Omega_i)]^2$, $i = 1, 2$ and $\|\sqrt{2\mu}\nabla \boldsymbol{z}_0\|_0 \leq \tilde{C}_{\text{stab}} \|\sqrt{2\mu}\nabla \boldsymbol{z}\|_0$. This falls out of the constructive proof in [29] and the fact that $\boldsymbol{V}_h$ is stable with respect to the one-dimensional pressure subspace of Q_h being subdomain-wise constant. For both $d = 2$ and $d = 3$, we formally set $\boldsymbol{v}_h := \boldsymbol{u}_h + \alpha S_h \boldsymbol{z}$ and show (20) if α and q_h are selected properly.

Now, we first consider the case $d = 3$. Using $q_h := -p_h - 2/3\mu \operatorname{div} \boldsymbol{u}_h \in Q_h$, the definition (19b) gives

$$\begin{aligned}
B(\boldsymbol{u}_h, p_h; \boldsymbol{v}_h, q_h) &= a_{\text{dev}}(\boldsymbol{u}_h, \boldsymbol{u}_h) - \frac{1}{3}(\mu \operatorname{div} \boldsymbol{u}_h, \operatorname{div} \boldsymbol{u}_h)_0 + c(p_h, p_h) + \frac{2}{3}c(p_h, \mu \operatorname{div} \boldsymbol{u}_h) \\
&\quad + \alpha a_{\text{dev}}(\boldsymbol{u}_h, S_h \boldsymbol{z}) - \alpha \frac{1}{3}(\mu \operatorname{div} \boldsymbol{u}_h, \operatorname{div} S_h \boldsymbol{z})_0 - \alpha(\operatorname{div} S_h \boldsymbol{z}, p_h)_0 \\
&\quad + (\operatorname{div} \boldsymbol{u}_h, p_h)_0 + \frac{2}{3}(\mu \operatorname{div} \boldsymbol{u}_h, \operatorname{div} \boldsymbol{u}_h)_0 - (\operatorname{div} \boldsymbol{u}_h, p_h)_0 \\
&= a_{\text{dev}}(\boldsymbol{u}_h, \boldsymbol{u}_h) + \frac{1}{3}(\mu \operatorname{div} \boldsymbol{u}_h, \operatorname{div} \boldsymbol{u}_h)_0 + c(p_h, p_h) + \frac{2}{3}c(p_h, \mu \operatorname{div} \boldsymbol{u}_h) \\
&\quad + \alpha a_{\text{dev}}(\boldsymbol{u}_h, S_h \boldsymbol{z}) - \alpha(p_h + \frac{1}{3}\mu \operatorname{div} \boldsymbol{u}_h, \operatorname{div} S_h \boldsymbol{z})_0.
\end{aligned}$$

Exploiting the fact that S_h is subdomain-wise defined and H^1-stable, we find in terms of (21b) that $|S_h \boldsymbol{z}|_{\boldsymbol{V}} \leq C_2 \|p_h + 1/3\mu \operatorname{div} \boldsymbol{u}_h\|_Q$, where C_2 only depends on the H^1-stability constant of S_h and C_1. Using (21a), we get

$$\begin{aligned}
B(\boldsymbol{u}_h, p_h; \boldsymbol{v}_h, q_h) &\geq \frac{1}{2}|\boldsymbol{u}_h|_{\boldsymbol{V}}^2 + \frac{1}{3}\|\sqrt{\mu} \operatorname{div} \boldsymbol{u}_h\|_0^2 + \alpha(1 - \frac{1}{2}C_2^2\alpha)\|p_h + \frac{1}{3}\mu \operatorname{div} \boldsymbol{u}_h\|_Q^2 \\
&\quad + c(p_h, p_h) + \frac{2}{3}c(p_h, \mu \operatorname{div} \boldsymbol{u}_h) - \alpha(p_h + \frac{1}{3}\mu \operatorname{div} \boldsymbol{u}_h, \operatorname{div}(S_h \boldsymbol{z} - \boldsymbol{z}))_0.
\end{aligned}$$

Due to the linear character of S_h, we have $S_h \boldsymbol{z} = \boldsymbol{z}_h + S_h \boldsymbol{z}_0$ and $\boldsymbol{z} - S_h \boldsymbol{z} = \boldsymbol{z}_0 - S_h \boldsymbol{z}_0$, and thus we can make use of the assumption on the stabilization term. Setting $C_3 := C_1 C_{\text{stab}} \tilde{C}_{\text{stab}}^2/2$, a straightforward computation yields

$$\begin{aligned}
B(\boldsymbol{u}_h, p_h; \boldsymbol{v}_h, q_h) &\geq \frac{1}{2}|\boldsymbol{u}_h|_{\boldsymbol{V}}^2 + \frac{1}{3}\|\sqrt{\mu} \operatorname{div} \boldsymbol{u}_h\|_0^2 + \alpha(1 - \frac{1}{2}(C_2^2 + C_3^2)\alpha)\|p_h + \frac{1}{3}\mu \operatorname{div} \boldsymbol{u}_h\|_Q^2 \\
&\quad + c(p_h, p_h) + \frac{2}{3}c(p_h, \mu \operatorname{div} \boldsymbol{u}_h) - c(p_h + \frac{1}{3}\mu \operatorname{div} \boldsymbol{u}_h, p_h + \frac{1}{3}\mu \operatorname{div} \boldsymbol{u}_h) \\
&= \frac{1}{2}|\boldsymbol{u}_h|_{\boldsymbol{V}}^2 + \frac{1}{6}(\|2\mu \operatorname{div} \boldsymbol{u}_h\|_Q^2 - \frac{1}{6}c(2\mu \operatorname{div} \boldsymbol{u}_h, 2\mu \operatorname{div} \boldsymbol{u}_h)) \\
&\quad + \alpha(1 - \frac{1}{2}(C_2^2 + C_3^2)\alpha)\|p_h + \frac{1}{3}\mu \operatorname{div} \boldsymbol{u}_h\|_Q^2.
\end{aligned}$$

Under the continuity assumption on the bilinear form $c(\cdot,\cdot)$ and by selecting α small enough, we find

$$B(\boldsymbol{u}_h, p_h; \boldsymbol{v}_h, q_h) \geq c(|\boldsymbol{u}_h|_{\boldsymbol{V}}^2 + \|p_h\|_Q^2 + \|2\mu \operatorname{div} \boldsymbol{u}_h\|_Q^2) \geq c(\|\boldsymbol{u}_h\|_{\boldsymbol{V}}^2 + \|p_h\|_Q^2),$$

whereas $\|(\boldsymbol{v}_h, q_h)\| \leq c(\|\boldsymbol{u}_h\|_{\boldsymbol{V}} + \|p_h\|_Q + \|2\mu \operatorname{div} \boldsymbol{u}_h\|_Q) \leq c(\|\boldsymbol{u}_h\|_{\boldsymbol{V}} + \|p_h\|_Q)$.

The proof for $d = 2$ with the weaker assumption on γ_0, i.e., $\gamma_0 < \infty$, follows essentially the same lines. Here we instead use $q_h := -p_h - \beta\mu \operatorname{div} \boldsymbol{u}_h \in Q_h$ with $\beta > 0$ small enough and exploit the fact that $a_{\mathrm{tr}}(\cdot,\cdot) = a_{\mathrm{dev}}(\cdot,\cdot)$. □

Remark 3. *We point out that the upper bound for γ_0 in Theorem 1 depends strongly on d. For $d = 2$, only continuity of $c(\cdot,\cdot)$ with respect to the pressure norm is required.*

Remark 4. *The proof shows that if on Ω_1 a non-trivial Dirichlet boundary part is prescribed, the assumption $\gamma_0 < 6$ in 3D can be weakened since then $|\cdot|_{\boldsymbol{V}}$ and $\|\cdot\|_{\boldsymbol{V}}$ are equivalent and $\|2\mu \operatorname{div} \boldsymbol{v}\|_Q$ is bounded by $|\boldsymbol{v}|_{\boldsymbol{V}}$. Exploiting this equivalence, we find an upper bound for γ_0 being larger than 6 but depending on a domain-dependent Korn-type constant which is in general unknown.*

3.2 An interior-penalty stabilization and local postprocessing

For $c(\cdot,\cdot)$ there are different options, e.g., we could choose a form which is related to the schemes discussed in [21]:

$$c(p,q) := \gamma \sum_{i=1}^{2} \frac{1}{2\mu_i} \sum_{F \in \mathcal{F}_h(\Omega_i)} \frac{|T_1^F| \cdot |T_2^F|}{|T_1^F| + |T_2^F|} [\![p]\!]|_F [\![q]\!]|_F, \quad \gamma > 0 \tag{22}$$

where $\mathcal{F}_h(\Omega_i)$ is the set of interior facets of the isoviscous subdomain Ω_i (i.e., facets on $\partial\Omega_i$ are excluded), and T_1^F, $T_2^F \in \mathcal{T}_h(\Omega_i)$ such that $\partial T_1^F \cap T_2^F = \bar{F}$. We point out that we do not penalize a pressure jump across the interface Γ_{12}.

A straightforward computation shows that if $\gamma < 3/2$ for $d = 3$ then the continuity assumption of Theorem 1 is satisfied. Integration by part and the properties of the Scott–Zhang operator with a proper scaling argument guarantee the second assumption on $c(\cdot,\cdot)$.

Remark 5. *Let $\mathcal{T}_h$ be a triangulation which is obtained by uniformly refining a mesh $\mathcal{T}_{4h}$ twice according to the procedure described in [4]. Then, since the pair $\boldsymbol{V}_h \times Q_{4h}$ is uniformly inf-sup stable, we can exclude those sub-faces of $\mathcal{F}_h$, which result from a uniform refinement of the faces of the base mesh $\mathcal{T}_{4h}$. Consequently, the choice of $\boldsymbol{z}$ in the proof of Theorem 1 has to be adapted, and $\boldsymbol{z}_0$ has to be selected such that it is equal to zero on the boundaries of all elements in $\mathcal{T}_{4h}$.*

As already pointed out hierarchically refined grids which result in patch-wise uniform fine meshes are of interest for massively parallel large scale simulations.

Although stabilized P_1-P_0 elements perform well with respect to computational cost and accuracy, the stabilization results in a loss of strong mass conservation. This turns out to be a possible disadvantage in coupled multi-physics problems where the velocity enters into an advective term. However, using a simple flux correction, we can easily recover exact element-wise mass-conservation in a local post-process. For convenience of the reader, we shall outline it here for completeness: For each facet F, we define an oriented facet flux by

$$j_{F;T} := \boldsymbol{u}_h \cdot \boldsymbol{n}_T + \Delta j_{F;T}, \quad \Delta j_{F,T} := \frac{\gamma}{2\mu|F|}\left(\frac{|T|\cdot|T^F|}{|T|+|T^F|}\right)(p_h|_T - p_h|_{T^F}),$$

where $\boldsymbol{n}_T$ is the outer unit normal with respect to the element T and $\partial T \cap \partial T^F = F$. Then, definition (22) guarantees local mass conservation, i.e., $\sum_{F\subset T}\int_F j_{F;T}\ ds = 0$. For many coupled problems it is sufficient to have a weak mass balance, i.e., to have access to $\int_F j_{F;T}\ ds$. If this is not the case, we can also lift the post-processed facet flux onto a BDM element; cf. e.g. [5] for further reading. In the case of continuous pressure approximations, this type of local post-process is also possible [14].

3.3 Possible generalizations

So far we restricted ourselves to the stabilized $P_1 - P_0$ setting. However, as the proof of Theorem 1 shows, we did not directly exploit the fact that we had been working with a piecewise constant pressure and a conforming piecewise linear velocity. The characteristic feature for the proof of our considered pairing is that we have $\operatorname{div} \boldsymbol{V}_h \subset Q_h$. Thus with minor modifications, the proof can be generalized to higher order stabilized P_k-P_{k-1} pairings provided that a suitable stabilization is specified.

Let us also briefly discuss the extension to other mesh types: We assume that $\mathcal{T}_h$ is a family of uniformly refined conforming hexahedral (quadrilateral) meshes and that we use trilinear (bilinear) finite elements in 3D (2D) for the velocity. The main difference between the Q_1-P_0 and the P_1-P_0 setting is that for the former case we do not have $\operatorname{div} \boldsymbol{V}_h \subset Q_h$. As a consequence, we cannot directly apply the techniques of the proof of Theorem 1. However, it is possible to adapt similar techniques also for the case of piecewise d-linear elements. Let us restrict ourselves to affinely mapped hexahedral elements, i.e., each element of $\mathcal{T}_h$ is affinely equivalent to the reference element $\hat{T} := (-1,1)^d, d = 3$. The following lemma guarantees a local equivalence of semi-norms.

Lemma 3. *Let $\hat{\boldsymbol{v}}_h \in [Q_1(\hat{T})]^d$, $d = 3$, then it holds*

$$2\|\operatorname{dev}\operatorname{sym}\nabla\hat{\boldsymbol{v}}_h\|^2_{0;\hat{T}} \le \frac{5}{3}\left(2\|\operatorname{dev}\operatorname{sym}\nabla\hat{\boldsymbol{v}}_h\|^2_{0;\hat{T}} - \frac{1}{3}\|\operatorname{div}\hat{\boldsymbol{v}}_h - \Pi_0\operatorname{div}\hat{\boldsymbol{v}}_h\|^2_{0;\hat{T}}\right). \tag{23}$$

Proof. A straightforward computation gives the required upper bound by solving a 24-dimensional generalized symmetric eigenvalue problem. □

Using the affine equivalence, the upper bound (23), and replacing $\operatorname{div} \boldsymbol{u}_h$ by $\Pi_0 \operatorname{div} \boldsymbol{u}_h$, in the definition of q_h and $\boldsymbol{z}$ in the proof of Theorem 1, we can easily adapt it for the $Q_1 - P_0$ case.

4 Numerical examples

Let us demonstrate the theoretical considerations by some illustrative examples. We study different test scenarios including traction boundaries and large viscosity jumps as well as several discretizations. In large scale computations, the use of stabilized equal-order P_1 discretizations (e.g. Brezzi and Pitkranta, [7]) are quite popular. Although they are formally not covered by our theory, we also study these discretizations below and show by numerical examples that no stability problems occur. Furthermore, we use a higher order discretization, namely a fourth order Scott-Vogelius element, which guarantees strong mass conservation and thus is automatically stable within our new formulation. Most of the implementations used in the following sections are based on the FEniCS (v 1.5.0) finite element framework [24]. However, the results for the stabilized equal-order P_1 discretization are obtained by a memory-efficient matrix-free implementation in the hierarchical hybrid grids (HHG) framework [3, 2, 11, 12, 10].

4.1 Traction-type boundary conditions

In this simple example, we shall demonstrate the effect of the modified bilinear form for applying traction boundary conditions which are not natural to incorporate in a standard gradient-based form of the Stokes problem. For simplicity, we consider $\mu = 1$ everywhere inside a rectangular domain $\Omega = [0, 5] \times [-1, 1]$, and we define a partition of the boundary into an outflow part $\Gamma_{\text{out}} = 5 \times [-1, 1]$, an inflow part $\Gamma_{\text{in}} = 5 \times (-1, 1)$ and a no-slip boundary $\Gamma_{\text{ns}} = \partial\Omega \backslash \overline{\Gamma_{\text{in}} \cup \Gamma_{\text{out}}}$. Given $\boldsymbol{u}_{\text{in}}(x, y) = (1 - y^2, 0)^\top$, the strong problem consists of finding $(\mathbf{u}, p)$ such that

$$
\begin{aligned}
-\Delta \boldsymbol{u} + \nabla p &= \mathbf{0}, && \text{in } \Omega, && \text{(24a)}\\
\operatorname{div} \boldsymbol{u} &= 0, && \text{in } \Omega, && \text{(24b)}\\
\boldsymbol{u} &= \boldsymbol{u}_{\text{in}}, && \text{on } \Gamma_{\text{in}}, && \text{(24c)}\\
\boldsymbol{u} &= \mathbf{0}, && \text{on } \Gamma_{\text{ns}}, && \text{(24d)}\\
\widehat{\boldsymbol{\sigma}}(\boldsymbol{u}, p) \cdot \boldsymbol{n} &= 0, && \text{on } \Gamma_{\text{out}}. && \text{(24e)}
\end{aligned}
$$

For comparison, we consider $\widehat{\boldsymbol{\sigma}}(\boldsymbol{u}, p) = \nabla \boldsymbol{u} - pI$ for the natural homogeneous Neumann boundary conditions for the gradient-based variant of Stokes' problem (we use $a(\cdot, \cdot) = (\nabla \cdot, \nabla \cdot)$ in the weak form), and we choose $\widehat{\boldsymbol{\sigma}}(\boldsymbol{u}, p) = \boldsymbol{\sigma}(\boldsymbol{u}, p)$ for the case of traction-free boundary conditions (we use $a(\cdot, \cdot) = a_{\text{tr}}(\cdot, \cdot)$ in the weak form). As a reference solution, we also consider the same problem, using the standard strain-based weak formulation which is compatible with traction-free boundaries (we use $a(\cdot, \cdot) = a_{\text{sym}}(\cdot, \cdot)$ in the weak form). For all cases,

we discretize our problem using the fourth order Scott–Vogelius element ($P_4 - P_3^{\text{disc}}$). This element is uniformly inf-sup stable for $d = 2$ and yields strongly divergence-free velocity solutions [37]. The mesh is a uniform mesh with $50 \times 20 \times 2 = 2\,000$ triangular elements.

Let us firstly remark that the stiffness-matrix for the gradient-based formulation and the stencil-modification only differs at the degrees of freedom associated with the outflow boundary. Hence, in a practical implementation the cross-derivatives only have to be considered in this region for consistency. Secondly, we mention that in contrast to the approach discussed in [23], we obtain a symmetric problem with provable stability and convergence. We finally report that for the above example the number of non-zeros in the viscous block only increases by around 0.75% compared to the saddle-point problem with the gradient-based operator, while the corresponding block in the strain-based Stokes operator requires roughly 89% more non-zeros. This reduction in the number of nonzero stencil entries leads to both, lower memory requirements and significant savings in terms of floating point operations that are required for the operator application. The savings are even greater in a three-dimensional setting.

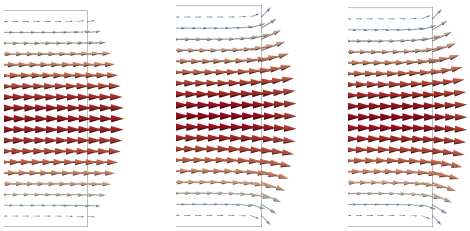

Figure 2: Outflow boundary: Plot of the velocity vectors in the outflow region for the gradient-based formulation (left), the stencil-modification (center) and the strain-based formulation (right).

In Fig. 2 we depict the outflow part of the domain for all three cases. We clearly observe that the boundary modification allows us to impose traction-free boundary conditions in a consistent manner also for a gradient-based formulation of the Stokes problem. Thus, the choice of the bilinear form does not dictate which boundary conditions shall be used in the whole outflow region. In particular, it is also possible to mix these two kinds of outflow conditions in a single code, giving more flexibility in modeling the flow-conditions that can be assumed past the outflow boundaries.

4.2 Interface problems

It is not straightforward to derive exact solutions for the three-dimensional interface problem and, in general, no closed-form solutions exist. Thus, we consider here a very simple benchmark problem extended from the two-dimensional case and more complex problems with numerically obtained reference solutions. If

not mentioned otherwise, we shall below use the stabilized $P_1 - P_0$ discretization of Section 3.1 together with the stabilization form (22) and $\gamma = 1$.

4.2.1 Couette flow

For the analytical benchmark, we consider the exact solution of Couette flow given in [18]: The domain is chosen as $\Omega = (0,1) \times (-\frac{1}{2}, \frac{1}{2}) \times (0,1)$ and the interface Γ is placed at $y = 0$ dividing Ω into $\Omega_1 = (0,1) \times (-\frac{1}{2}, 0) \times (0,1)$ and $\Omega_2 = (0,1) \times (0, \frac{1}{2}) \times (0,1)$. We then set $\boldsymbol{f} = (3\mu, 0, 0)^\top$ and boundary conditions such that the exact solution is given by

$$\boldsymbol{u} = (\tfrac{1}{2}(1 - x^2), xy, 0)^\top, \qquad p(x,y) = 2x\mu - \tfrac{1}{2}(\mu_1 + \mu_2),$$

where $\mu_1 = 1$ and $\mu_2 =\in \{10^{-3}\}$. Our initial mesh is an interface-resolving uniform tetrahedral mesh with mesh-size $h = \frac{1}{4}$, and we study the convergence rates under uniform refinement, see Table 1. As our numerical studies indicate, the expected optimal convergence rates are obtained.

level	$\|\boldsymbol{u} - \boldsymbol{u}_h\|_0$	rate	$\|\boldsymbol{u} - \boldsymbol{u}_h\|_{\mathbf{V}}$	rate	$\|p - p_h\|_Q$	rate
0	9.5263e-03	–	1.1065e-01	–	6.4989e-01	–
1	2.4913e-03	1.93	5.3663e-02	1.04	1.3437e-01	2.27
2	6.4259e-04	1.95	2.6567e-02	1.01	3.9015e-02	1.78
3	1.6373e-04	1.97	1.3256e-02	1.00	1.5737e-02	1.30
4	4.1456e-05	1.98	6.6258e-03	1.00	7.3621e-03	1.09

Table 1: Results for the non-isoviscous Couette flow example: $\mu_1/\mu_2 = 10^3$

4.2.2 Viscosity layers

In this test, we consider a more complex problem for which we have no analytical solution available. We again use the unit cube $\Omega := (0,1)^3$ and define four subdomains $\Omega_i := (0,1)^2 \times (i/4, (i+1)/4)$, $i = 0, \dots 3$, where we set the viscosities in each of these subdomains to $\mu_i = 10^i$, i.e., each additional layer increases the viscosity by an order of magnitude. For the boundary condition, we apply free-slip conditions, and we choose the forcing as $\mathbf{f} = (0, 0, -\cos(2\pi x)\cos(2\pi y)\sin(\pi z))^\top$. We choose an initial discretization of the domain into 4^3 cubes which are then each subdivided into six tetrahedra to form the triangulation at the lowest level.

We compare the discrete solutions $(\boldsymbol{U}_h, P_h)$ obtained for the strain-based Stokes problem with cross-derivatives (i.e., using $a_{\text{sym}}(\cdot,\cdot)$ in the weak formulation) to the discrete solutions $(\boldsymbol{u}_h, p_h)$ obtained using the modified form without global cross-couplings using $a_{\text{tr}}(\cdot,\cdot)$ in the weak formulation. Moreover, to demonstrate the effect of the modified bilinear form, we also compare the strain-based formulation with the gradient-based version that is obtained by using $a(\cdot,\cdot) = (\mu\nabla\cdot, \nabla\cdot)$ in the weak formulation. It should be noted that we are comparing numerical solutions here. However, since the strain-based Stokes

discretization is a widely-accepted model for non-isoviscous incompressible flow, we are confident that the reference solution suffices to demonstrate that our method converges optimally in more general situations.

level	$\frac{\|\boldsymbol{U}_h-\boldsymbol{u}_h\|_0}{\|\boldsymbol{U}_h\|_0}$	rate	$\frac{\|\boldsymbol{U}_h-\boldsymbol{u}_h\|_{\boldsymbol{V}}}{\|\boldsymbol{U}_h\|_{\boldsymbol{V}}}$	rate	$\frac{\|P_h-p_h\|_Q}{\|P_h\|_Q}$	rate
0	2.9339e-01	–	3.0574e-01	–	1.7826e-01	–
1	3.0164e-01	–	4.1250e-01	–	2.2806e-01	–
2	3.1359e-01	–	4.8968e-01	–	2.5337e-01	–
3	3.1095e-01	–	5.1625e-01	–	2.6207e-01	–
4	3.0922e-01	–	5.2354e-01	–	2.6458e-01	–

Table 2: Layer example using stabilized $P_1 - P_0$ elements: relative errors obtained using the gradient-based vs. the strain-based Stokes operator on a series of uniformly refined meshes (negative or small rates are not displayed).

To obtain a better understanding about the magnitude of the errors involved, we divide the discrete error norms by the respective norm of the reference solution. The results listed in Tables 2 and 3 show that for this seemingly simple setup, the gradient-based form produces results which are up to 50% off in the energy norm and up to 30% off in the L^2-norm of the velocity.

level	$\frac{\|\boldsymbol{U}_h-\boldsymbol{u}_h\|_0}{\|\boldsymbol{U}_h\|_0}$	rate	$\frac{\|\boldsymbol{U}_h-\boldsymbol{u}_h\|_{\boldsymbol{V}}}{\|\boldsymbol{U}_h\|_{\boldsymbol{V}}}$	rate	$\frac{\|P_h-p_h\|_Q}{\|P_h\|_Q}$	rate
0	1.7799e-01	–	1.6976e-01	–	1.8562e-01	–
1	9.9843e-02	0.83	1.0584e-01	0.68	1.3775e-01	0.43
2	3.6872e-02	1.43	5.2343e-02	1.01	5.9058e-02	1.22
3	1.0722e-02	1.78	2.2601e-02	1.21	2.1634e-02	1.44
4	2.8306e-03	1.92	1.0224e-02	1.14	8.3074e-03	1.38

Table 3: Layer example using stabilized $P_1 - P_0$ elements: relative errors obtained using the stencil-modified vs. the strain-based Stokes operator.

The consistently modified form however leads to an equally convergent scheme as the strain-based reference implementation with cross-derivatives. For comparison, we depict the pressure fields obtained using all approaches in Figure 3.

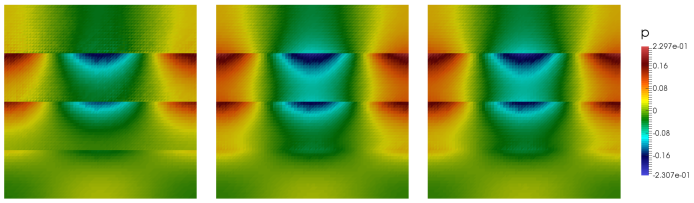

Figure 3: Viscosity layers: Plot of the pressure field ($y = 0.5$) on the finest mesh level for the gradient-based formulation (left), the stencil modification (center) and the reference solution obtained by the strain-based Stokes formulation (right).

4.2.3 Viscosity columns

Next, we consider the same domain, mesh-sequence, boundary conditions and forcing terms as in the previous example of Section 4.2.2, but we choose a different viscosity field. We define the columns $\Omega_2 = (0, \frac{1}{4})^2 \times (0,1) \cup (\frac{3}{4}, 1)^2 \times (0,1)$ and denote the remaining domain by $\Omega_1 := \Omega \backslash \overline{\Omega_2}$. Furthermore, we set $\mu_1 = 1$ in Ω_1 and $\mu_2 = 10$ in Ω_2. We note that for this setup, we expect a singularity in the pressure at the edge where the normal is discontinuous, see Figure 4 for some illustration.

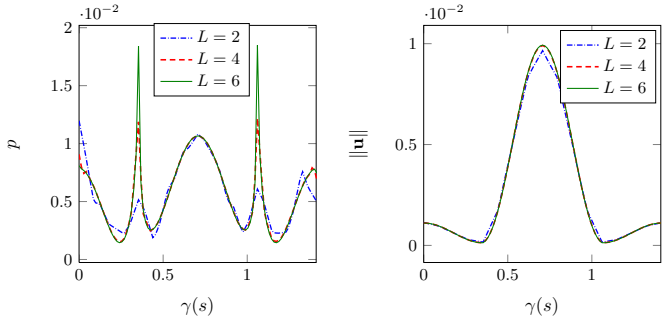

Figure 4: Viscosity columns: Mesh study of the discrete pressure (left), magnitude of the velocity (right) on a 1D line $\gamma(s)$ cutting diagonally through the domain.

In this example, we first show the errors between the consistent stencil formulation and the classical strain based formulation on the same mesh for the theoretically covered $P_1 - P_0$ case and for the popular $P_1 - P_1$ discretization which is not covered by our theory. Note that we use a sparse-matrix implementation in the $P_1 - P_0$ case, whereas for the $P_1 - P_1$ discretization a memory-efficient matrix-free implementation is used, for which we compute finer mesh levels, i.e., up to level 6 in Tables 5, 7 and 8.

level	$\frac{\|\boldsymbol{U}_h - \boldsymbol{u}_h\|_0}{\|\boldsymbol{U}_h\|_0}$	rate	$\frac{\|\boldsymbol{U}_h - \boldsymbol{u}_h\|_{\boldsymbol{V}}}{\|\boldsymbol{U}_h\|_{\boldsymbol{V}}}$	rate	$\frac{\|P_h - p_h\|_Q}{\|P_h\|_Q}$	rate
0	6.6276e-02	–	7.1282e-02	–	9.2322e-02	–
1	3.2402e-02	1.03	5.1998e-02	0.45	7.5894e-02	0.28
2	2.1200e-02	0.61	4.3765e-02	0.24	6.6193e-02	0.19
3	1.9921e-02	0.08	4.2069e-02	0.05	6.5935e-02	0.00
4	1.9985e-02	0.00	4.1805e-02	0.00	6.6048e-02	0.00

Table 4: Column example using stabilized $P_1 - P_0$ elements: relative errors obtained using the gradient-based vs. the strain-based Stokes operator.

Tables 4 and 5 show the results for the inconsistent gradient based formulation. We again observe that, independently of the selected discretization, the gradient-based form of the Stokes problem converges to a different solution than the strain-based approach. In this example, the lack of consistency introduces

a relative L^2 error of roughly 2% which for both discretizations under consideration does not vanish under mesh refinement.

level	$\frac{\|\boldsymbol{U}_h - \boldsymbol{u}_h\|_0}{\|\boldsymbol{U}_h\|_0}$	rate	$\frac{\|\boldsymbol{U}_h - \boldsymbol{u}_h\|_{\boldsymbol{V}}}{\|\boldsymbol{U}_h\|_{\boldsymbol{V}}}$	rate	$\frac{\|P_h - p_h\|_Q}{\|P_h\|_Q}$	rate
2	7.4550e-02	–	9.2581e-02	–	2.7609e01	–
3	3.0104e-02	1.36	6.0398e-02	0.64	1.8003e-01	0.64
4	2.1227e-02	0.52	4.7364e-02	0.36	1.1179e-01	0.70
5	2.0227e-02	0.07	4.3246e-02	0.13	8.0546e-02	0.48
6	2.0098e-02	0.01	4.2121e-02	0.04	6.9997e-02	0.20

Table 5: Column example using stabilized $P_1 - P_1$ elements: relative errors obtained using the gradient-based vs. the strain-based Stokes operator.

The results for the consistently modified scheme are listed in Tables 6 and 7. As is theoretically expected, the stencil based formulation eliminates the consistency error in both cases and achieves good rates of convergence in the respective norms.

level	$\frac{\|\boldsymbol{U}_h - \boldsymbol{u}_h\|_0}{\|\boldsymbol{U}_h\|_0}$	rate	$\frac{\|\boldsymbol{U}_h - \boldsymbol{u}_h\|_{\boldsymbol{V}}}{\|\boldsymbol{U}_h\|_{\boldsymbol{V}}}$	rate	$\frac{\|P_h - p_h\|_Q}{\|P_h\|_Q}$	rate
0	1.1023e-01	–	1.1881e-01	–	1.4978e-01	–
1	5.0381e-02	1.12	7.1554e-02	0.73	1.0439e-01	0.52
2	1.7661e-02	1.51	3.4368e-02	1.05	5.0572e-02	1.04
3	4.9194e-03	1.84	1.5852e-02	1.11	2.5225e-02	1.00
4	1.2678e-03	1.95	7.5513e-03	1.06	1.2397e-02	1.02

Table 6: Column example using stabilized $P_1 - P_0$ elements: relative errors obtained using the stencil-modified vs. the strain-based Stokes operator.

level	$\frac{\|\boldsymbol{U}_h - \boldsymbol{u}_h\|_0}{\|\boldsymbol{U}_h\|_0}$	rate	$\frac{\|\boldsymbol{U}_h - \boldsymbol{u}_h\|_{\boldsymbol{V}}}{\|\boldsymbol{U}_h\|_{\boldsymbol{V}}}$	rate	$\frac{\|P_h - p_h\|_Q}{\|P_h\|_Q}$	rate
2	1.3890e-02	–	2.4189e-02	–	7.7949e-02	–
3	4.2414e-03	1.79	1.2312e-02	1.02	4.8195e-02	0.72
4	1.1370e-03	1.94	6.2303e-03	1.00	2.5287e-02	0.95
5	2.9449e-04	1.97	3.1918e-03	0.98	1.3274e-02	0.94
6	7.5985e-05	1.96	1.6675e-03	0.94	6.8175e-03	0.97

Table 7: Column example using stabilized $P_1 - P_1$ elements: relative errors obtained using the stencil-modified vs. the strain-based Stokes operator.

Here, we not only compare the errors on the same mesh levels, but also with respect to a solution obtained on two subsequent refinements of the maximum considered mesh level. Table 8 shows the results for the stabilized $P_1 - P_1$ discretization. We observe a slightly reduced rate for the L^2 error of the velocity and the pressure-error, which may be due to the impact of the singularity on the asymptotic order of convergence. It shall also be noted that the difference between the two consistent formulations on the same mesh is roughly by a factor of five smaller than the discretization error that can be estimated with help of the high fidelity solution.

level	$\frac{\|\boldsymbol{U}_h-\boldsymbol{u}_h\|_0}{\|\boldsymbol{U}_h\|_0}$	rate	$\frac{\|\boldsymbol{U}_h-\boldsymbol{u}_h\|_{\boldsymbol{V}}}{\|\boldsymbol{U}_h\|_{\boldsymbol{V}}}$	rate	$\frac{\|P_h-p_h\|_Q}{\|P_h\|_Q}$	rate
2	1.1073e-01	–	4.8731e-01	–	6.9077e-01	–
3	2.0170e-02	2.57	1.3751e-01	1.90	1.5605e-01	2.24
4	5.7436e-03	1.85	6.9474e-02	1.01	8.2428e-02	0.94
5	1.5953e-03	1.86	3.4779e-02	1.01	4.3231e-02	0.94
6	4.2928e-04	1.90	1.7104e-02	1.03	2.2758e-02	0.93

Table 8: Column example using stabilized $P_1 - P_1$ elements: relative errors obtained using the stencil-modified vs. the strain-based Stokes operator. For the reference we take a level 8 solution of the strain-based formulation.

5 Massively parallel multigrid

In this section, we study the influence of the specific operator formulation on the performance and scalability of an all-at-once multigrid method. Our solver considers the discrete saddle point problem

$$\begin{pmatrix} A & B^\top \\ B & -C \end{pmatrix} \begin{pmatrix} \boldsymbol{u}_h \\ p_h \end{pmatrix} = \begin{pmatrix} \boldsymbol{f}_h \\ 0 \end{pmatrix},$$

where the submatrices are associated with nodal basis functions and the associated bilinear forms of the equal-order scheme. Let us recall that only A depends on the specific operator formulation, while C depends on the stabilization that is employed, and B denotes the weak divergence operator. We use a variable multigrid V_{var}-cycle with three pre- and post-smoothing steps of Uzawa-type, i.e., in the $(k+1)$th-smoothing step we solve

$$\begin{aligned} \mathbf{u}_{k+1} &= \mathbf{u}_k + \hat{A}^{-1}(\mathbf{f} - A\mathbf{u}_k - B^\top \mathbf{p}_k), \\ \mathbf{p}_{k+1} &= \mathbf{p}_k + \hat{S}^{-1}(B\mathbf{u}_{k+1} - C\mathbf{p}_k), \end{aligned} \tag{25}$$

where $\hat{A}$ and $\hat{S}$ denote suitable approximations for A and the Schur-complement S, respectively; cf. [41, 35, 10]. In particular for $\hat{A}$, we use a symmetric hybrid parallel variant of a row-wise red-black colored Gauss-Seidel method and for $\hat{S}$, we select the forward variant of the Gauss-Seidel method applied to the stabilization matrix C, with $\omega = 0.3$ as under-relaxation factor. Standard restriction and prolongation operators define the transfer between the meshes.

Before comparing the formulations in the context of the all-at-once multigrid method, let us comment on the coarse grid solver. As long as a direct coarse grid solver is used, the use of $a_{\text{tr}}(\cdot,\cdot)$ is not problematic. However, quite often, especially in case of moderate viscosity contrasts, preconditioned Krylov subspace solvers may give shorter run-times. Here we apply a block-diagonal preconditioned minimal residual (MINRES) method, where a Jacobi-preconditioned CG-iteration is used for the velocity block and a lumped mass-matrix for the pressure. The stopping criteria are selected such that the multigrid convergence rates do not deteriorate.

However, this preconditioner cannot be straightforwardly applied for $a_{\text{tr}}(\cdot,\cdot)$. Therefore, we exchange the stencil-formulation by the strain-based Stokes for-

mulation on the coarse grid and apply the preconditioned MINRES method there. Although the stencil-formulation is more efficient and faster than the strain-based Stokes formulation when iterating on finer grids, the coarse grid performance is dominated by communication. Thus, we can choose the more expensive but positive definite strain-based Stokes formulation on the coarse mesh.

A naive replacement of $a_{\mathrm{tr}}(\cdot,\cdot)$ by $a_{\mathrm{sym}}(\cdot,\cdot)$ on the coarse level will fail. We recall that $a_{\mathrm{tr}}(\cdot,\cdot)$ is only equivalent to $a_{\mathrm{sym}}(\cdot,\cdot)$ when $\mathrm{div}(\boldsymbol{u}) = 0$. On the coarse grid, the right hand side of the discrete mass conservative equation is typically non-zero. Thus, the two discrete coarse level systems

$$\begin{pmatrix} A_{\mathrm{tr}} & B^\top \\ B & -C \end{pmatrix} \begin{pmatrix} \boldsymbol{z} \\ \boldsymbol{q} \end{pmatrix} = \begin{pmatrix} \boldsymbol{r}_1 \\ \boldsymbol{r}_2 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} A_{\mathrm{sym}} & B^\top \\ B & -C \end{pmatrix} \begin{pmatrix} \boldsymbol{z} \\ \boldsymbol{q} \end{pmatrix} = \begin{pmatrix} \boldsymbol{r}_1 \\ \boldsymbol{r}_2 \end{pmatrix}$$

do not approximate the same physical system. Here all matrices are associated with the coarse level, and the subindex in A reflects the selected bilinear form. To account for this coarse grid inconsistency between the two systems, we modify the coarse grid right-hand side of the momentum equation. Replacing $\boldsymbol{r}_1$ by $\boldsymbol{r}_1 - B^\top M^{-1} \boldsymbol{r}_2$, where M is the lumped mass matrix, i.e., a diagonal matrix, we find that

$$\begin{pmatrix} A_{\mathrm{tr}} & B^\top \\ B & -C \end{pmatrix} \begin{pmatrix} \boldsymbol{z} \\ \boldsymbol{q} \end{pmatrix} = \begin{pmatrix} \boldsymbol{r}_1 \\ \boldsymbol{r}_2 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} A_{\mathrm{sym}} & B^\top \\ B & -C \end{pmatrix} \begin{pmatrix} \boldsymbol{z} \\ \boldsymbol{q} \end{pmatrix} = \begin{pmatrix} \boldsymbol{r}_1 - B^\top M^{-1} \boldsymbol{r}_2 \\ \boldsymbol{r}_2 \end{pmatrix}$$

are approximations of the same physical system, and we can use the modified right hand side in combination with A_{sym} for the coarse mesh solver.

To test the multigrid performance, we use two different architectures and base our implementation on the massively-parallel geometric multigrid framework HHG [3, 2, 11, 12, 10]. Firstly, we consider a two-socket server system equipped with Intel Xeon E5-2699 processors, 32 cores in total and 256 GB main memory. Due to the matrix-free aproach even on such a low-cost machine, we can solve for 10^9 DoF. Secondly, the IBM Blue Gene/Q supercomputing system JUQUEEN[1] with a peak performance of more than 5.9 petaflop/s, currently ranked on position 11 of the TOP500 list (46th edition, November 2015) is used. On this system, Stokes flow problems with 10^{13} DoF can be considered [10]. As test example, we reuse the model problem of Section 4.2.3.

In Table 9, we consider the refinement levels $L = 5, 6, 7, 8$ on the intel machine and level $L = 2$ with 19 652 DoF for the coarse grid. The asymptotic convergence rate ρ for the residual reduction is approximated by using the convergence rate after 200 iterations with rescaling in each step. For all our settings, the convergence rates are robust with respect to the problem size and are similar valued. Both, the modified stencil and the gradient-based formulations, show almost constant rates $\rho = 0.241$ over the levels and are additional slightly better than the ones of the strain-based Stokes operator $\rho = 0.255$ on $L = 8$.

Increasing the refinement level, the run-times for the modified stencil and gradient-based form are up to a factor 1.73 faster than the strain-based Stokes

[1] www.fz-juelich.de

	strain-based form		modified stencil form		gradient-based form	
level	T [s]	ρ	T [s]	ρ	T [s]	ρ
5	1.18	0.296	1.12	0.244	1.34	0.240
6	3.02	0.280	2.45	0.242	2.64	0.240
7	13.95	0.263	9.04	0.241	9.26	0.241
8	93.74	0.255	54.08	0.241	54.22	0.241

Table 9: Time of one $V_{\text{var}}(3,3)$-cycle and asymptotic convergence rates for the different formulations on the Intel machine

operator. The modified stencil and the strain-based Stokes-operator need a similar number of coarse grid iterations to reach the stopping criteria, while the pure gradient-based operator requires more iterations such that the run-times are slightly higher than for the Stencil operator.

On JUQUEEN, we perform a weak scalability study and report the total time and the time without coarse grid solver (w.c.). For all three formulations, the run-times without coarse grid time stay almost constant in the weak scaling. Including the coarse grid, the run-times increase due to the suboptimal coarse grid solver. In the weak scaling, the difference between the strain-based Stokes and modified stencil operator is between 7.65s and 8.78s. When neglecting the coarse grid, the difference increases to a factor of around 1.61 between the strain-based Stokes operator and the modified stencil and the gradient-based operator.

		strain-based form		modified stencil form		gradient-based form	
proc.	DoF	T [s]	T (w.c.) [s]	T [s]	T (w.c.) [s]	T [s]	T (w.c.) [s]
128	$5.4 \cdot 10^8$	25.83	21.92	18.18	14.31	18.61	14.28
768	$4.3 \cdot 10^9$	32.04	21.93	23.26	13.90	25.44	13.90
8 192	$3.4 \cdot 10^{10}$	39.37	21.33	31.50	13.69	36.86	13.27
65 536	$2.8 \cdot 10^{11}$	71.09	21.41	63.23	13.27	69.22	13.26

Table 10: Weak scaling of one $V_{\text{var}}(3,3)$-cycle and asymptotic convergence rates for the different formulations on JUQUEEN. All timings are given in seconds.

6 Geophysical example

In our last example, we study the influence of the different formulations on the solution quality for a simplified geophysical setup. It is widely assumed that there is a huge viscosity contrast between the lower Earth's mantle and the asthenosphere, which is a mechanically weak layer in the uppermost mantle; cf., e.g., [20, 25, 19, 30]. The precise depth of the asthenosphere is still unknown, although it is accepted that depth and viscosity contrast are closely coupled. In order to take the Haskell constraint into account and to study the influence of two different asthenospheric depths $d_a = 660\,[\text{km}]$ and $d_a = 410\,[\text{km}]$, we choose

the piecewise constant viscosity

$$\nu = \begin{cases} 10^{22}\,[\mathrm{Pa\,s}] & \text{for } r < |\mathbf{x}| < R - d_a, \\ 10^{21}\left(\frac{d_a}{1000[km]}\right)^3\,[\mathrm{Pa\,s}] & \text{for } |\mathbf{x}| \geq R - d_a, \end{cases} \tag{26}$$

see Weismüller et al. [39], resulting in a viscosity jump of 34.78 and 145.09, respectively. Here, R is the radius of the Earth and $R - r$ is the thickness of the mantle close to 3000 [km]. The right hand side of the incompressible Stokes system is given by a temperature driven buoyancy term $(-\rho\mathbf{g})$, i.e. the product of a density field ρ [kg/m^3] and the gravitational acceleration within the mantle $\mathbf{g}$ [m/s^2], which we prescribe as a vector of magnitude 10 [m/s^2] pointing towards the center of the Earth. The density ρ is obtained from a tomographic model of seismic wave speeds within the Earth [15], converted to densities with the mineralogical model of Stixrude et al. [38]; see also [1]. At the outer boundary, Dirichlet conditions for the velocity are specified according to tectonic plate data [26], and at the inner boundary free-slip conditions are assumed. Both types of boundary conditions yield $\boldsymbol{u} \cdot \boldsymbol{n} = 0$.

For the rather large viscosity contrasts, which are considered here, our coarse mesh solver consumes a considerable amount of the computation time. Thus, we shall not compare performance figures here as we did in the previous section, but only examine the impact on the solution quality. An investigation of other coarse grid solvers, combining algebraic and geometric multigrid techniques, can be found for instance in the recent contribution of Rudi et al. [34].

We perform our geophysical simulations using the Uzawa-multigrid V_{var}–cycles introduced in Section 5 to solve the geophysical problem with the different formulations. The initial mesh consists of an icosahedral mesh of the spherical shell with 922, 560 tetrahedrons. This mesh, which resolves the viscosity jump, is refined twice to obtain the coarse level mesh and then four more times to build up the multilevel hierarchy. Altogether, this results in a computational problem consisting of $2.5 \cdot 10^9$ degrees of freedom. The maximal velocities for each formulation and the differences to the strain-based formulation, which we again take as a reference, are reported in Table 11. The difference in the numbers are a first indicator that the gradient-based form yields non-physical results for this setup in the sense that these are not consistent with the strain-based model. In both settings, the maximal difference between our new formulation and the classical strain-based formulation is roughly 1% whereas the gradient-based model results in a maximal difference of more than 10%.

In the top row of Figure 5, we present simulation results for the strain-based operator and asthenosphere depths $d_a = 660$ (left) and $d_a = 410$ (right). We scale the color bar to a maximum velocity of 15 [m/s] and observe that the highest velocities occur in the thin asthenospheric layer. The smaller d_a, the higher the maximal velocity is. For both d_a, a similar structure in the velocity distribution is obtained. To get a better feeling for the different formulations and the effect of the physical inconsistent interface coupling in case of the gradient-

	maximal velocity			maximal difference	
	reference	mod. stencil	grad.-based	mod. stencil	grad.-based
$d_a = 660$	17.67	17.67	19.02	0.18	3.51
$d_a = 410$	28.13	28.12	30.10	0.17	3.24

Table 11: Maximal velocity magnitudes (in [m/s]) for the different formulations and maximal absolute differences (in [m/s]) for asthenosphere depths $d_a = 660$ (left) and $d_a = 410$.

based formulation, we also provide a scaled relative point-wise error plot

$$\mathrm{Error}_k(x) = \frac{\|\boldsymbol{u}_\mathrm{e}(x) - \boldsymbol{u}_k(x)\|}{\max(\|\boldsymbol{u}_\mathrm{e}(x)\|, 5 \cdot 10^{-2} \|\boldsymbol{u}_\mathrm{e}\|_{L^\infty(\Omega)})}, \tag{27}$$

with $k \in \{\mathrm{g}, \mathrm{s}\}$, denoting the gradient-based (g) and the stencil-based (s) discrete solution, respectively. The coloring of $\mathrm{Error}_\mathrm{g}$ is scaled to 50% and of $\mathrm{Error}_\mathrm{s}$ to 2.5%. In the center row, $\mathrm{Error}_\mathrm{s}$ is shown, and $\mathrm{Error}_\mathrm{g}$ is depicted in the third row. For both d_a, the gradient-based form shows high differences triggered by the incorrect physical formulation at the interface and the inner boundary. We observe that this is a global effect and pollutes into the interior of the lower mantle. However, for the new stencil based formulation no inconsistency error enters at the interface, and thus the relative difference is quite small. For the larger asthenospheric depth, there is almost no visible difference while for the smaller layer moderate differences can be observed in regions where we have small velocities.

Acknowledgements

This work was supported (in part) by the German Research Foundation (DFG) through the Priority Programme 1648 "Software for Exascale Computing" (SPPEXA). The authors gratefully acknowledge the Gauss Centre for Supercomputing (GCS) for providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS share of the supercomputer JUQUEEN at Jülich Supercomputing Centre (JSC).

References

[1] S. Bauer, H.-P. Bunge, D. Drzisga, B. Gmeiner, M. Huber, L. John, M. Mohr, U. Rüde, H. Stengel, C. Waluga, J. Weismüller, G. Wellein, M. Wittmann, and B. Wohlmuth. Solution Techniques for the Stokes System: A priori and a posteriori modifications, resilient algorithms. In *SPPEXA proceedings*, 2015. accepted.

[2] B. Bergen, T. Gradl, U. Rüde, and F. Hülsemann. A massively parallel multigrid method for finite elements. *Computing in Science and Engineering*, 8(6):56–62, 2006.

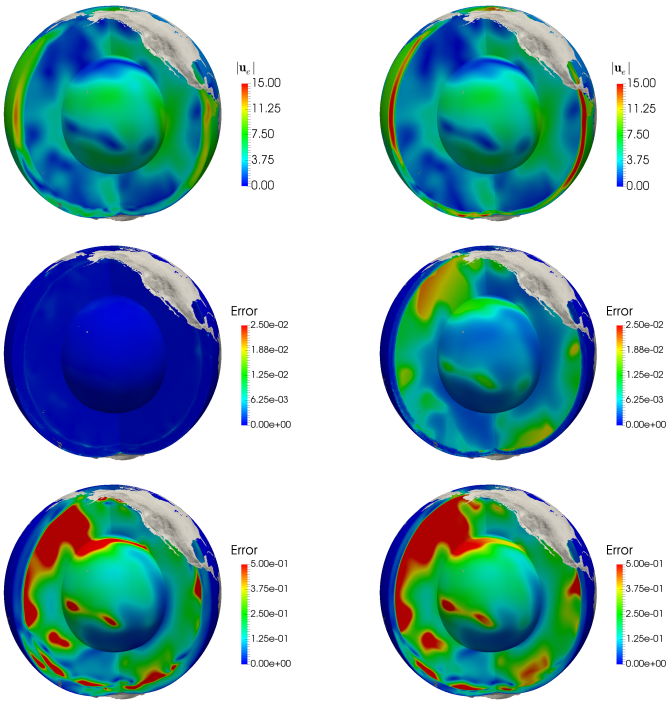

Figure 5: Study of different asthenosphere depths $d_a = 660$ (left) and $d_a = 410$ (right): Velocity-magnitude of the strain-based Stokes formulation (8) (top). Relative error for the new stencil modification (11) (middle) and the gradient-based formulation (bottom).

[3] B. Bergen and F. Hülsemann. Hierarchical hybrid grids: data structures and core algorithms for multigrid. *Numerical Linear Algebra with Applications*, 11:279–291, 2004.

[4] J. Bey. Tetrahedral grid refinement. *Computing*, 55(4):355–378, 1995.

[5] D. Boffi, F. Brezzi, and M. Fortin. *Mixed finite element methods and applications*. Springer, 2013.

[6] D. Boffi, N. Cavallini, F. Gardini, and L. Gastaldi. Local mass conservation of Stokes finite elements. *Journal of Scientific Computing*, 52(2):383–400, 2012.

[7] F. Brezzi and J. Pitkäranta. On the stabilization of finite element approximations of the Stokes equations. In W. Hackbusch, editor, *Efficient Solutions of Elliptic Systems*. Springer, 1984.

[8] B. R. Cousins, S. Le Borne, A. Linke, L. G. Rebholz, and Z. Wang. Efficient linear solvers for incompressible flow simulations using Scott–Vogelius finite elements. *Numerical Methods for Partial Differential Equations*, 29(4):1217–1237, 2013.

[9] T. Gerya. *Introduction to Numerical Geodynamic Modelling*. Cambridge University Press, 2009.

[10] B. Gmeiner, M. Huber, L. John, U. Rüde, and B. Wohlmuth. A quantitative performance analysis for Stokes solvers at the extreme scale, 2015. submitted, arXiv:1511.02134.

[11] B. Gmeiner, H. Köstler, M. Stürmer, and U. Rüde. Parallel multigrid on hierarchical hybrid grids: a performance study on current high performance computing clusters. *Concurrency and Computation: Practice and Experience*, 2012.

[12] B. Gmeiner, U. Rüde, H. Stengel, C. Waluga, and B. Wohlmuth. Performance and Scalability of Hierarchical Hybrid Multigrid Solvers for Stokes Systems. *SIAM J. Sci. Comput.*, 37(2):C143–C168, 2015.

[13] B. Gmeiner, U. Rüde, H. Stengel, C. Waluga, and B. Wohlmuth. Towards textbook efficiency for parallel multigrid. *Numer. Math. Theory Methods Appl.*, 8, 2015.

[14] B. Gmeiner, C. Waluga, and B. Wohlmuth. Local mass-corrections for continuous pressure approximations of incompressible flow. *SIAM Journal on Numerical Analysis*, 52(6):2931–2956, 2014.

[15] S. P. Grand, R. D. van der Hilst, and S. Widiyantoro. Global seismic tomography: A snapshot of convection in the earth. *GSA Today*, 7:1–7, 1997.

[16] S. Gross and A. Reusken. *Numerical methods for two-phase incompressible flows*, volume 40. Springer Science & Business Media, 2011.

[17] B. Hager and R. Clayton. *Mantle Convection*, chapter Constraints on the structure of mantle convection using seismic observations, flow models, and the geoid, pages 657–764. Gordon and and Breach, New York, 1989.

[18] P. Hansbo, M. G. Larson, and S. Zahedi. A cut finite element method for a Stokes interface problem. *ArXiv e-prints*, May 2012.

[19] R. Hartley, G. Roberts, N. White, and C. Richardson. Transient convective uplift of an ancient buried landscape. *Nature Geosci.*, 4:562–565, 2011.

[20] N. A. Haskell. The motion of a fluid under a surface load. *Physics*, 6:265–269, 1935.

[21] T. J. R. Hughes and L. P. Franca. A new finite element formulation for computational fluid dynamics: VII. the Stokes problem with various well-posed boundary conditions: symmetric formulations that converge for all velocity/pressure spaces. *Computer Methods in Applied Mechanics and Engineering*, 65(1):85–96, 1987.

[22] K. Ito and Z. Li. Interface conditions for stokes equations with a discontinuous viscosity and surface sources. *Applied mathematics letters*, 19(3):229–234, 2006.

[23] A. Limache, S. Idelsohn, R. Rossi, and E. Oñate. The violation of objectivity in Laplace formulations of the Navier–Stokes equations. *International journal for numerical methods in fluids*, 54(6-8):639–664, 2007.

[24] A. Logg, K.-A. Mardal, and G. N. Wells. *DOLFIN: a C++/Python Finite Element Library*, volume 84 of *Lecture Notes in Comp. Science and Engineering*, chapter 10. Springer, 2012.

[25] J. X. Mitrovica. Haskell [1935] revisited. *J. Geophys. Res.*, 101:555–569, 1996.

[26] R. D. Müller, M. Sdrolias, C. Gaina, and W. R. Roest. Age, spreading rates, and spreading asymmetry of the world's ocean crust. *Geochem. Geophy. Geosy.*, 9:1525–2027, 2008.

[27] P. Neff, D. Pauly, and K.-J. Witsch. Poincaré meets Korn via Maxwell: Extending Korn's first inequality to incompatible tensor fields. *Journal of Differential Equations*, 258(4):1267 – 1302, 2015.

[28] M. Neilan. Discrete and conforming smooth de Rham complexes in three dimensions. *Mathematics of Computation*, 2015.

[29] M. A. Olshanskii and A. Reusken. Analysis of a Stokes interface problem. *Numerische Mathematik*, 103(1):129–149, 2006.

[30] R. Parnell-Turner, N. White, T. Henstock, B. Murton, J. Maclennan, and S. M. Jones. A continuous 55 million year record of transient mantle plume activity beneath Iceland. *Nature Geosci.*, 7:914–919, 2014.

[31] O. Pironneau. *Finite element methods for fluids*. Wiley Chichester, 1989.

[32] W. Pompe. Counterexamples to Korn's inequality with non-constant rotation coefficients. *Math. Mech. Solids*, 16:172–176, 2011.

[33] A. Quarteroni and A. Valli. *Domain Decomposition Methods for Partial Differential Equations*. Numerical mathematics and scientific computation. Clarendon Press, 1999.

[34] J. Rudi, A. Malossi, T. Isaac, G. Stadler, M. Gurnis, P. Staar, Y. Ineichen, C. Bekas, A. Curioni, and O. Ghattas. An extreme-scale implicit solver for complex PDEs: Highly heterogeneous flow in Earth's mantle. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 5:1–5:12, 2015.

[35] J. Schöberl and W. Zulehner. On Schwarz-type smoothers for saddle point problems. *Numer. Math.*, 95(2):377–399, 2003.

[36] L. Scott and S. Zhang. Finite element interpolation of nonsmooth functions satisfying boundary conditions. *Mathematics of Computation*, 54(190):483–493, 1990.

[37] L. R. Scott and M. Vogelius. Norm estimates for a maximal right inverse of the divergence operator in spaces of piecewise polynomials. *Modélisation mathématique et analyse numérique*, 19(1):111–143, 1985.

[38] L. Stixrude and C. Lithgow-Bertelloni. Thermodynamics of mantle minerals – I. Physical properties. *Geophys. J. Int.*, 162:610–632, 2005.

[39] J. Weismüller, B. Gmeiner, S. Ghelichkhan, M. Huber, L. John, B. Wohlmuth, U. Rüde, and H. P. Bunge. Fast asthenosphere motion in high-resolution global mantle flow models. *Geophys. Res. Lett.*, 42(18):7429–7435, 2015.

[40] S. Zhang. A new family of stable mixed finite elements for the 3d Stokes equations. *Mathematics of Computation*, 74(250):543–554, 2005.

[41] W. Zulehner. Analysis of iterative methods for saddle point problems: a unified approach. *Math. Comp.*, 71(238):479–505, 2002.